

ATTORNEY DOCKET NO.  
062986.0214  
1151.00

PATENT APPLICATION

1

SYSTEM AND METHOD FOR ORDERING OF DATA TRANSFERRED OVER  
MULTIPLE CHANNELS

TECHNICAL FIELD OF THE INVENTION

The present invention relates in general to computer information processing and more particularly to a system and method for ordering data transferred over multiple channels.

5

BACKGROUND OF THE INVENTION

Conventional data transfer schemes may provide flow control for a single channel. However, these schemes do not address the issue of splitting a transfer over multiple physical or virtual channels. In environments where the ordering of the data is important, splitting the data transfer over multiple channels may cause the data to sent in an order different from originally desired. For example, graphics systems may require the processing of data in a certain order to effectively generate an image for display that reflects the intended scene to be viewed. The destination will not be able to effectively process the data since it is received in a different order than when it was generated. Therefore, it is desirable to have a capability to re-order the data at the destination despite the use of multiple channels for data transfer.

SUMMARY OF THE INVENTION

From the foregoing, it may be appreciated by those skilled in the art that a need has arisen for a technique that can re-order data at a destination after the data has been transferred over multiple channels while providing flow control for the data transfer. In accordance with the present invention, a system and method for ordering data transferred over multiple channels are provided that substantially eliminate or greatly reduce disadvantages and problems associated with conventional data transfer techniques.

According to an embodiment of the present invention, there is provided a method for ordering data transferred over multiple channels that includes generating a data packet in response to a flow control credit. A sequence number is generated and inserted into the data packet. One of a plurality of channels to transfer the data packet is selected and the data packet is transferred over the selected one of the plurality of channels

The present invention provides various technical advantages over conventional data transfer techniques. For example, one technical advantage is to associate a sequence number to data packets so that they can be placed into the proper sequence at the destination. Another technical advantage is to place data packets received from a source in an appropriate sequence despite transferring the data packets over multiple channels. Yet another technical advantage is to integrate re-ordering of data packets with flow control of the data packet transfer. Other examples may be readily ascertainable by those skilled in the art from the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and the advantages thereof, reference is now made to the following description taken in conjunction with the accompanying drawings, wherein like reference numerals represent like parts, in which:

FIGURE 1 illustrates a block diagram of a multiple channel data transfer system;

FIGURE 2 illustrates a block diagram of a source for data packet generation in the multiple channel data transfer system;

FIGURE 3 illustrates a block diagram of a sequence number unit in the source;

FIGURE 4 illustrates a block diagram of a destination for processing data packets in the multiple channel data transfer system.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram of a multiple channel data transfer system 10. System 10 includes a source 12 that generates data packets for transfer over multiple request channels 14. Data packets are transferred over multiple request channels 14 through a network 16 to a destination 18. Destination 18 processes the data packets received over multiple request channels 14 and provides appropriate response packets to source 12 over multiple response channels 20 through network 16. In an exemplary scenario, system 10 may be a graphics system wherein source 12 is a node with one or more processors generating data packets and destination 18 is a graphics processor that generates images in response to the data packets. Multiple request channels 14 are needed as the amount of data generated by a processor of the node may exceed the bandwidth of a single channel. Network 16 may be any data transfer medium that may include one or more of a local area network, wide area network, Ethernet network, Internet, or any conventional network for transferring data.

Each data packet generated at source 12 may be placed onto any available one of the multiple request channels 14. For each data packet, source 12 determines an available channel according to capacity and flow control instructions. Examples of flow control implementations in single channel systems can be found in co-pending U.S. Application Serial No. 09/409,805 and U.S. Patent No. 6,154,794, which are each hereby incorporated by reference herein. In a multi-channel implementation, a first packet may be sent on a first channel and a second packet may be sent on another

channel. Because data packets may be sent on different channels, destination 18 may receive and buffer the second packet before the first packet even though the first packet was generated prior to generation of the second packet. Since the processing of data packets in the correct order is important, destination 18 needs to recognize that the first packet is to be processed before the second packet.

FIGURE 2 shows a block diagram for the generation of data packets at source 12. Source 12 includes a source request channel controller 22, a sequence number unit 24, a response channel controller 26, and a credit counter unit 28. Source request channel controller 22 generates data packets from data received from a data source, such as a processor, in source 12 and selects among multiple request channels for the transferring of the data packets. Alternatively, the data source may provide the data in packet form for transfer by source request channel controller 22.

Source request channel controller 22 inserts a sequence number as provided by sequence number unit 24 for each generated data packet. The sequence number is used by destination 18 in order to re-order the received data packets into the proper sequence. Upon inserting a sequence number into a data packet, source request channel controller 22 transfers the data packet onto a selected one of multiple request channels 14. Source request channel controller 22 issues an increment command to sequence number unit 24 so that the current sequence number provided by sequence number unit 24 to source request channel controller 22 is advanced to the next number in the sequence. The sequence number may be reset

by the data source for each block of data to be transferred.

Source request channel controller 22 also issues a decrement command to credit counter unit 28 to indicate that a credit was used in transferring the data packet over the selected one of the multiple request channels 14. Credit counter unit 28 is initialized with a number of credits by a reset signal from the data source. The number of credits is decremented after transfer of each data packet from source request channel controller 22. Credit counter unit 28 provides a flow control indication to the data source to control the amount of data provided by the data source. As long as the number of credits is above a desired threshold value, the flow control indication informs the data source that it may continue to provide data. If the number of credits falls below the desired threshold value, the flow control indication informs the data source to cease providing data until the desired threshold value is again exceeded. The data source may force an interrupt to do a context switch in order that other tasks may be performed in the event that the flow control indication maintains the data cessation status for a certain period of time. The data source will then return to the transfer of data packets. The number of credits is incremented in response to response channel controller 26 receiving a reply from destination 18 over one of multiple response channels 20. The reply includes a credit indicating that destination 18 has processed a data packet and has capability to receive another data packet.

FIGURE 3 shows a block diagram of sequence number unit 24. Sequence number unit 24 includes a current

sequence number register 30, a maximum sequence number register 32, comparator 34, increment unit 36, increment selectors 38 and 40, reset logic 42, and increment logic 44. Upon initialization, a reset signal from the data source causes a sequence number value of zero to be loaded into current sequence number register 30. Increment selector 38 receives the value in current sequence number register 30 and an incremented sequence number value from increment unit 36. Increment counter 38 provides the value from current sequence number register 30 to increment selector 40 unless an increment signal has been received from source request channel controller 22 indicating that a data packet with the value in current sequence number register 30 has been transferred. Increment selector 40 provides the value in current sequence number register 30 back thereto unless a reset signal is received from the data source or an increment signal is received from source request channel controller 22. The reset signal causes the zero value to be loaded into current sequence number register 30. The increment signal causes an incremented value to flow through increment selectors 38 and 40 for loading into current sequence number register 30.

Current sequence number register 30 may also be loaded with the zero value according to a maximum value in maximum sequence number register 32. The value in current sequence number register 30 is compared to the value in maximum sequence number register 32 at comparator 34. If these values are equal, comparator 34, through increment logic 44 and reset logic 42, causes the zero value at increment selector 40 to be loaded into current sequence number register 30. In this manner, a



large data packet group may be accommodated without loss of packet ordering since sequence numbers will be freed up and available for re-use once a data packet has been processed. Current sequence number register 30 and maximum sequence number register 32 may be programmable to provide a range of sequence numbers to simplify buffering at destination 18. Also, the programmable nature of these registers allows for varying the size of buffers receiving data packets at destination 18.

FIGURE 4 shows the processing performed at destination 18. Destination 18 includes a destination request channel controller 50, a re-order buffer write port controller 52, a re-order buffer 53, a re-order buffer valid bit unit 54, a re-order buffer read port controller 56, a credit response unit 57, and a destination response channel controller 58. In operation, destination request channel controller 50 receives data packets from source 12 over multiple request channels 14. Destination request channel controller 50 provides the data packets to a re-order buffer write port controller 52. Ordering of data packets is not guaranteed over multiple request channels 14. To regain the proper ordering of the data packets, re-order buffer write port controller uses the sequence numbers in the data packets to directly index into re-order buffer 53 through its write port. Re-order buffer valid bit unit 54 maintains a valid bit for each portion of re-order buffer 53. When re-order buffer write port controller 52 causes data packets to be stored in re-order buffer 53, a corresponding valid bit is set. Re-order buffer read port controller 56 monitors the valid bits. When a next sequential valid bit is set, the data

packets associated with the immediately preceding valid bit in the corresponding portion of re-order buffer 53 are read from its read port and provided to a data destination. Once the data packets have been read from re-order buffer 53, the corresponding valid bit is cleared. All valid bits may be cleared in response to a reset signal. Through use of the sequence numbers, data packets are read from re-order buffer 53 in their original order of generation. As packets are read from re-order buffer 53, credit response unit 57 returns flow control credits back to source 12 through destination response channel controller 58 over multiple response channels 20.

Thus, it is apparent that there has been provided, in accordance with the present invention, a system and method for ordering data transferred over multiple channels that satisfies the advantages set forth above. Although the present invention has been described in detail, it should be understood that various changes, substitutions, and alterations may be made herein. For example, the functionality for providing sequence numbers and flow control credits may be performed in any number of ways other than the implementation described above for appropriate sequence number and credit processing. Other examples may be readily ascertainable by those skilled in the art and may be made herein without departing from the spirit and scope of the present invention as defined by the following claims.